Met Office

Comparing satellite altimetry with ocean models of the North-West Shelf

Forecasting Research
Technical Report No: 581

July 19, 2013

Robert R. King and Matthew Martin

Abstract

The Forecasting Ocean Assimilation Model (FOAM) is the short-range operational ocean prediction system at the Met Office. Although many data-types are assimilated into deep ocean configurations of FOAM, currently only SST is assimilated in the shelf seas and, in general, there has been little work on assimilating other data-types in the shelf seas. This is due to a number of issues including complex vertical co-ordinate systems, the mismatch between the cadence of observations and actual shelf dynamics, and the need to filter altimeter products for variations caused by atmospheric pressure and wind speed.

The TAPAS project has produced a new experimental sea level anomaly (SLA) product without along-track filtering or sub-sampling, which details the corrections applied to the sea surface height (SSH) measurements. Here we present a comparison of the predicted SSH fields from the Atlantic Margin Model (AMM) implementation of FOAM with the TAPAS observations. These comparisons have shown that the best match is found when neither of the atmospheric corrections provided with TAPAS are applied to the observations resulting in a $\sim$10 cm standard deviation in the observation-model difference. A power spectrum analysis indicates that while the high frequency mesoscale signal is not well-resolved by the observations or model, the surge component could be usefully assimilated. A further comparison of the TAPAS observations with the more mature CS3X surge model found that, in general, the CS3X surge model residuals are only 70% of those for FOAM.

1 Introduction

Continental shelf seas are of great physical and economic importance: the majority of the ocean's tidal energy is dissipated in the shelf seas [Egbert and Ray, 2000], they can be highly biologically productive, and are host to oil and gas extraction platforms, shipping and fisheries. Current ocean prediction systems for the shelf seas provide operational forecasts of hydrodynamic variables and predict the well-mixed and seasonal-stratification patterns and fronts present in such shallow seas [e.g., O'Dea et al., 2012]. The output of these models can be used, among other things, to understand the propagation of sonar waves using the 3D density structure, to drive ecosystem models, and are used by search and rescue and the oil and gas industry [De Mey et al., 2009].

The Met Office's short-range operational ocean forecasting system, the Forecasting Ocean Assimilation Model (FOAM), is run in various configurations to produce daily analyses and 7-day forecasts of the ocean state, including temperature, salinity, currents, sea-ice concentration and sea surface height. For global ocean configurations, the NEMOVAR assimilation scheme allows the assimilation of temperature and salinity profiles, in situ and satellite sea surface temperature (SST) observations, sea-ice concentration, and satellite altimeter observations of the sea surface height (SSH). Such data assimilation has proved to be an important part of operational ocean forecasting [Martin et al., 2007, Cummings et al., 2009].

In the shelf seas however, only SST observations are currently assimilated into FOAM. This is partly due to the scales of the dynamics of the dominant physical processes acting on the shelf

© Crown Copyright 2013

compared to the relative spatial and temporal sparseness of the available observations. In particular, the tidal interaction with the shelf can lead to sea level changes of several metres over a few hours, and storm surges can give rise to sea level changes of up to 2–3 metres over a 6 hour period [Flowerdew et al., 2010]. This contrasts with altimeter SSH observations with tracks typically a few 100 km apart and repeat cycles of 10–30 days and, on the other hand, high cadence tide gauge data which can be used to validate model predictions, but are present only around coastlines.

Generally, data assimilation in the shelf seas is more challenging than in the open ocean. In addition to the mismatch of spatial and temporal scales, the shelf seas can be stratified in the summer months and well-mixed in winter when the solar illumination is not sufficient to induce stratification. Since the assimilation of SSH in the open ocean employs a technique whereby the position of the thermocline is altered, assimilation of SSH where there is no stratification (and hence no thermocline) must be by another method. However, the assimilation of SSH observations could improve ocean forecasts as they provide a vertically integrated measure of the temperature and salinity of a water column, information on barotropic effects (including the tides and surge) and can be used to infer the geostrophic component of surface currents. Their assimilation would therefore be complementary to SST observations, which are limited to the ocean surface, but can be combined with profile observations to estimate the 3-dimensional ocean density structure. Such an estimate is hindered by the relative lack of profile data, although data from gliders can address this to some extent, but altimeter observations, on the other hand, provide better spatial coverage of the entire domain.

In this report we compare predicted SSH fields from the Atlantic Margin Model (AMM) implementation of FOAM with a new Aviso/CLS Sea Level Anomaly (SLA) product to investigate the possibility of assimilating these altimeter observations into FOAM. To also investigate the current accuracy of surge prediction in FOAM in the North-West Shelf (NWS), we have performed a similar comparison between the TAPAS altimeter observations and the CS3X surge model. Although the TAPAS products are currently only available in a small number of specific regions, subsequent improvements to the ocean model and data assimilation scheme in the NWS will be similarly applicable in other shelf seas such as the Persian Gulf if the data become available. Indeed, Dobricic et al. [2012] have shown an improvement in the SLA analyses from the Mediterranean Forecasting System [Pinardi and Coppini, 2010] when the TAPAS data are used in this less tidally-affected region.

The first part of this report details the available altimeter observations, supplied as part of the TAPAS project, and goes on to describe the FOAM model with which the observations are compared. The results of the comparison are presented in Section 4 and a further comparison of the TAPAS observations with the CS3X surge model is presented in Section 5. Finally, these comparisons are discussed in Section 6 and our conclusions and recommendations for the assimilation of these data are given in Section 7.

© Crown Copyright 2013

2 Observations

The observations used here are an experimental data-set produced by CLS/Aviso as part of the TAPAS (Tailored Altimetry Products for Assimilation Systems) project and include observations for January 2008 to December 2009 from three satellites (ENVISAT, Jason-1 and Jason-2). These delayed time, along-track SLA observations are provided without along-track sub-sampling or filtering and without the usual corrections for high-frequency signals. The unfiltered and non-subsampled data have a sampling frequency of 1Hz, providing a horizontal resolution of $\sim$7 km, and the SLA measurement is constructed from the measured sea surface height (SSH) and a mean dynamic topography [Rio et al., 2011].

For every observation the associated corrections are provided, allowing judgements on the need for, and efficacy of, each correction. The along-track resolution of the altimeter observations ($\sim$7 km) closely matches the model spatial resolution. However, the altimeter tracks are necessarily sparsely sampled in time due to the repeat cycles of $\sim$10 (Jason-1/2) and $\sim$35 days (ENVISAT) as demonstrated by Fig. 1.

Although the TAPAS data are currently only available for the European area, the purpose of the project is to determine whether such data-sets can be assimilated into operational forecasting systems. If this proves worthwhile, similar data-sets may become available for more regions and improvements propagated to operational forecasting in other shelf seas.

3 Ocean Model

The Atlantic Margin Model (AMM) configuration of FOAM uses NEMO v3.2 [Madec, 1999] with a time-step of 150 s and includes atmospheric pressure forcing and tides. It covers the North-West European Shelf and part of the North-East Atlantic ocean with a resolution of 1/15° latitude by 1/9° longitude, giving a horizontal resolution of $\sim$7 km [see Storkey et al., 2010, O'Dea et al., 2012]. The AMM domain extends from 40°S, 20°W to 65°N, 13°E.

This domain contains regions of shallow water, a continental shelf and regions of deep ocean (depth $\gtrsim$ 5,000 m) with a complex bathymetry (e.g., the Norwegian trench). The resulting dynamics, due to tidal mixing and the seasonal variation in wind strength and solar irradiation, leads to seasonally stratified and well-mixed regions with marked boundaries. Wind forcing and atmospheric pressure can produce storm surges, and the internal ocean dynamics contain mesoscale eddies and fronts.

Since the model output provides the sea surface height, while the observations are supplied as sea level anomalies, from each model time-step we subtracted a 1-year mean SSH field (as a proxy for a mean dynamic topography) to produce a model SLA. Examples of instantaneous and mean model SSH fields are shown in Fig. 1. To compare the observations and model values, for each observed point the model values were linearly interpolated first in space (at several time-steps

© Crown Copyright 2013

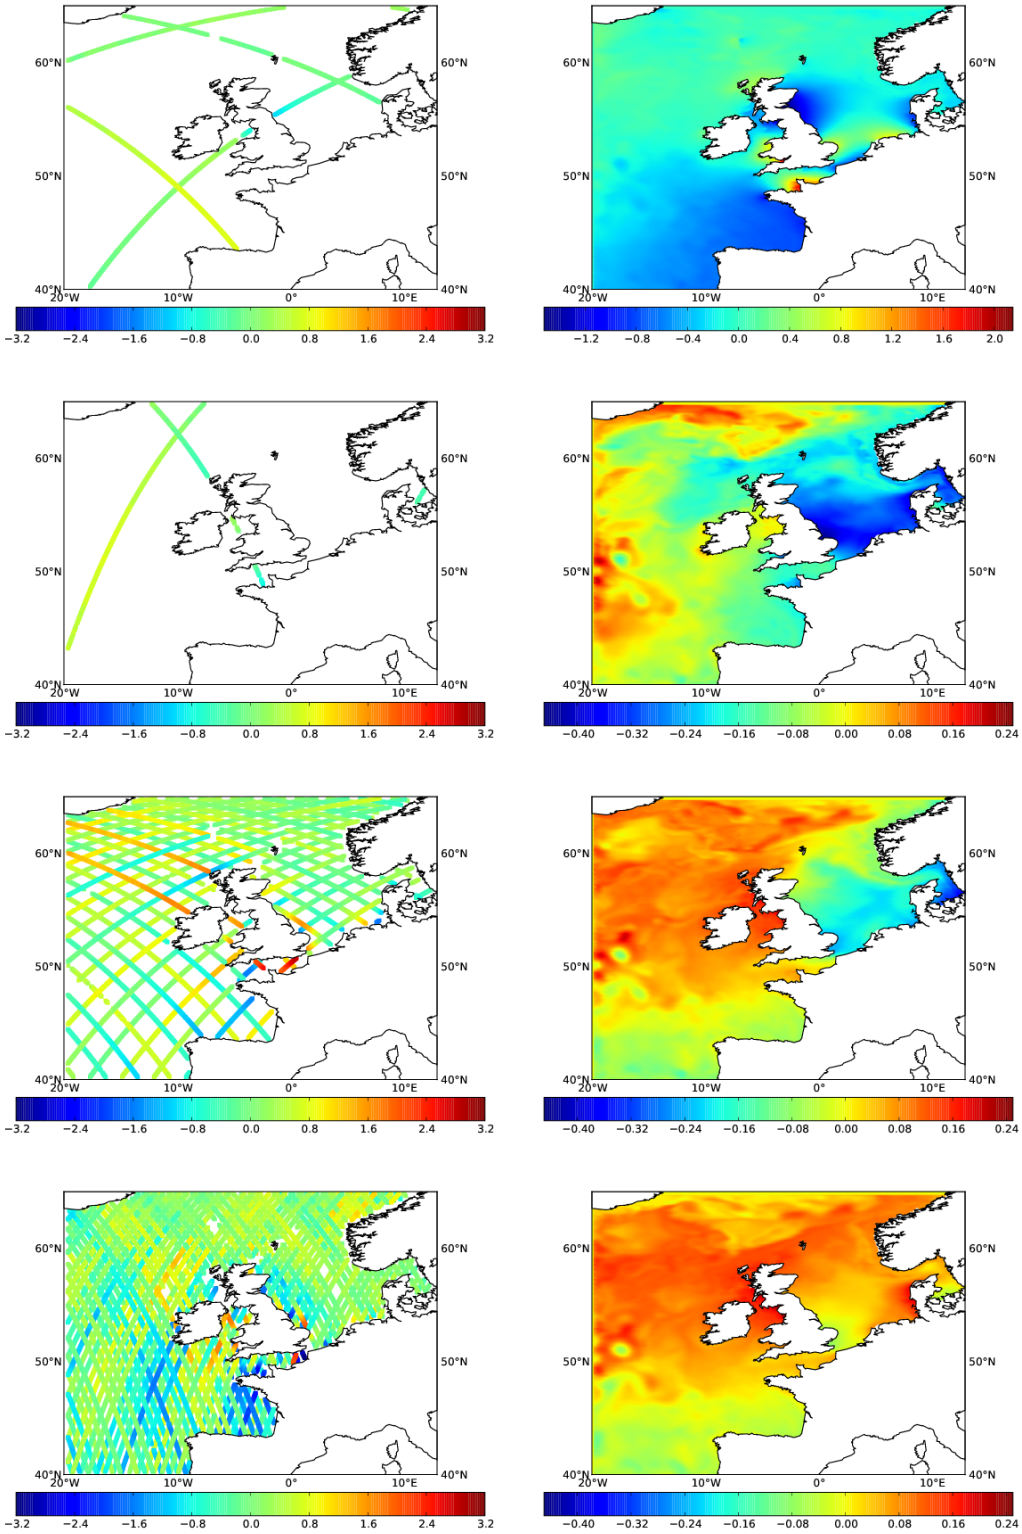

Figure 1: A comparison of the spatial coverage of satellite altimeter observations (left) with the model SLA fields (right) for January 2008 with the colour representing the observed or simulated sea level anomaly (SLA) in metres. The left column shows Jason-1 (top) and ENVISAT (2nd top) tracks over a single day and over the repeat period for each satellite (10 and 35 days, respectively). The right column shows an instantaneous model SLA field on the same date as the Jason-1 tracks (top), the mean SLA field on that same day, and the mean fields over the respective repeat periods.

© Crown Copyright 2013

Correction	Description
MOG2D	2D Gravity Wave Model Used to correct ocean response to atmospheric pressure and wind forcing.
LFIB	Low-Frequency Inverse Barometer Used to correct ocean response to atmospheric pressure forcing.
TIDE	Ocean Tide Used to remove tidal signature.
LWE	Long Wavelength Error Used to remove effect of inaccurate satellite orbit determination.

Table 1: The corrections applied to the satellite altimetry measurements.

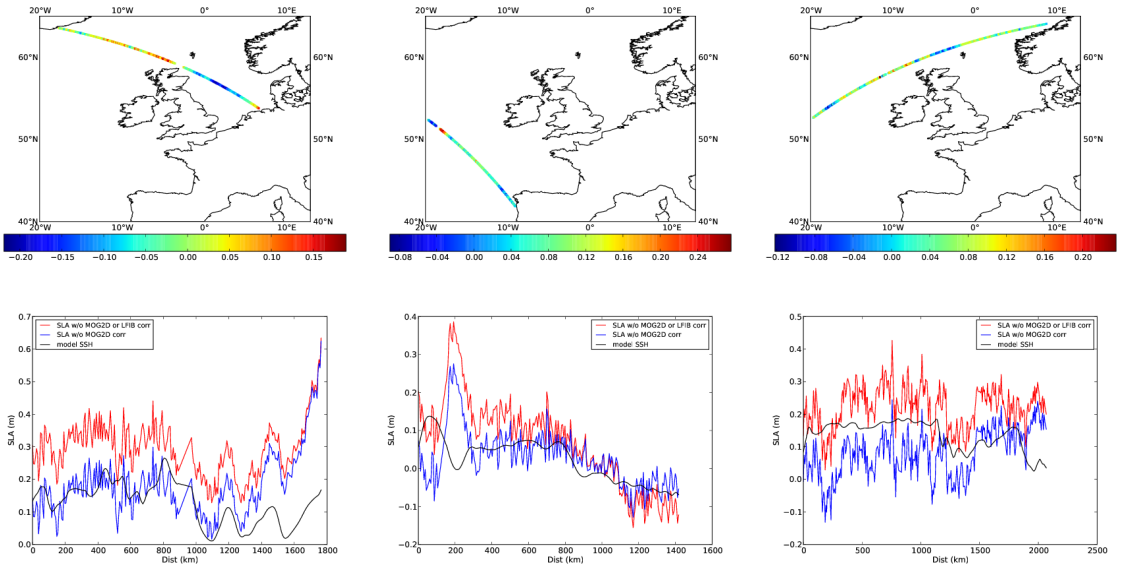

Figure 2: Three example Jason-1 tracks showing the track positions (top row, colour indicates the fully-corrected SLA) and a comparison of the model SLA to the observations (bottom row) without the MOG2D correction (blue) and without the MOG2D or LFIB corrections (red). Tides have been removed from the model (black) using a Doodson filter (see Sect. 4.3).

bracketing the observations) and then in time.

Our available model output spanned 2007–2008, but the first available TAPAS observations are from January 2008. As such we have used ENVISAT observations over a full year (01 Jan. 2008 – 31 Dec. 2008) and Jason-1 observations over the 10.5 months (01 Jan. 2008–19 Oct. 2008) prior to the realignment of the satellite's orbit. The large difference between the model resolution and the spatial and temporal satellite coverage of the domain is in evident in Fig. 1 where the instantaneous model field is dominated by tides (note the colour-scale of ± 2.5 m), but the mean model fields over a single day and over the repeat cycle of each satellite show smaller scale variations due to the lower amplitude dynamics in the region.

© Crown Copyright 2013

4 NEMO/TAPAS comparison

4.1 Choice of observational corrections

As discussed in Sec. 2, the sea surface height observations supplied as part of the TAPAS project are provided with a set of corrections, which are summarised in Table 1. These four corrections aim to remove the systematic bias introduced by an inaccurate satellite orbit determination, the effects of wind and atmospheric pressure forcing, and the large tidal signature present.

The low-frequency inverse barometer effect (LFIB) correction is a static formulation intended to remove the response of the ocean to atmospheric pressure forcing. It is produced using linearly interpolated 6-hourly ECMWF atmospheric pressure fields. The complementary MOG2D correction, intended to remove higher frequency effects due to atmospheric pressure and wind forcing, is produced using a 2-dimensional barotropic model forced by the pressure and wind speeds from the same ECMWF analysis as for the LFIB correction.

The ocean tide correction supplied with the observations is produced from the output of a global ocean tide model [Ray, 1999]. Since this is a relatively low resolution ($\sim$0.5°) *global* tide model, it may not reliably reproduce all tidal constituents on the shelf. Finally, the long wavelength error (LWE) correction is a small term intended to correct the imperfectly known orbit of each satellite.

Similar corrections are routinely applied to altimeter products assimilated into deep ocean operational forecasting systems. These corrections have also been used to allow the production of a consistent, long-term record of sea level changes. However, this introduces inconsistencies with shelf-seas models which include the effect of wind and/or pressure forcing from associated NWP models. We must therefore decide which of these corrections to apply to the TAPAS SLA measurements before comparing to our ocean model output.

The long wavelength error (LWE) is a small correction affecting the measurement, but unconnected to ocean physics and should therefore be applied. The tidal correction should only be applied where the model either does not include tidal effects or the model output has been filtered to suppress the tidal signature. Finally, the LFIB and MOG2D corrections are both intended to remove from the observation the signal due to atmospheric effects on the sea level, i.e. the forcing due to changes in atmospheric pressure and wind speeds. The FOAM-AMM model includes the effect of forcing due to the atmospheric pressure and wind, and so to provide the best match between the observations and the physical processes captured by the model, the altimetry measurements should not have these corrections applied. Fig. 2 shows three example tracks, where the observations have had the tidal and LWE corrections applied, compared with the interpolated model output.

4.2 Spatial map of temporal variations

To compare the effects of applying or excluding these corrections, we have produced maps of the absolute mean difference and standard deviation of the difference between the observations and

© Crown Copyright 2013

model fields averaged over a year of observations (only 10.5 months were used for Jason-1 as its orbit was altered toward the end of 2008). Since the amplitude of the measured signal is often dominated by tides, we have employed the tidal corrections supplied with the observations and for the model fields we have filtered the tidal signal using a Doodson filter, which uses a weighted combination of the signal at 39 times centred on the time of interest.

Figures 3 and 4 show the comparison of Jason-1 and ENVISAT observations with four sets of corrections: i) all corrections except MOG2D are applied, ii) all except MOG2D and LFIB are applied, iii) all except LFIB are applied, and iv) all corrections are applied. From these maps it is clear that the absolute mean and standard deviation of the observation-model differences are typically higher in the North Sea and around coastlines than in the western parts of the domain. This is especially pronounced when all corrections are applied (lower panels of Figs. 3 and 4).

The absolute mean difference and standard deviation of the differences, taken over the full year and over the full domain, are marginally lower for the case where the MOG2D correction is not applied than for the case where neither the MOG2D nor LFIB corrections are applied. However, an inspection of the spatial plots of these statistics shows that, when neither the MOG2D nor LFIB corrections are applied, the residuals in the North Sea are reduced. This same effect is seen on a monthly time-scale. These monthly changes across the full domain are demonstrated by the Taylor diagrams in Figs. 5 and 6 which show large month-to-month differences.

Furthermore, when we compare the mean observation-model differences (rather than using the absolute mean) with neither of the atmospheric corrections applied, a North-South gradient is apparent (see Fig. 7). This could be related to the construction of the SLA. The SLA observations supplied by Aviso/CLS provide the sea surface height with respect to a mean sea surface height determined from several years of observations. However, here we have used a model mean over a single year to produce our model SLA. In Fig. 7 we show the difference between the 1-year model mean field and the Rio et al. [2011] MDT, which is broadly consistent with the North-South gradient seen in the mean observation-model differences. With a longer time series of model fields, this disparity may be reduced. On the other hand, the disparity may be a result of the construction of the mean observed SSH field if this was produced from the SLA with the atmospheric corrections applied. However, the same conclusion on the choice of corrections is reached, and similar standard deviations are found, when the Rio et al. MDT is used instead of our 1-year model mean.

From the discussion in the previous section, we should expect the best match between the observations and model to be when the atmospheric forcing corrections (i.e. MOG2D and LFIB) are not applied to the observations. The comparisons above broadly support this, but it is not immediately clear that the LFIB correction should be excluded. Figures 5 and 6 suggest a lower seasonal variation in the observation-model differences when only the MOG2D correction is excluded. However, it appears that the exclusion of both the MOG2D and LFIB corrections leads to the best match across the domain (especially in the North Sea) and so we will not apply either of these atmospheric corrections to the satellite observations.

© Crown Copyright 2013

The complicated effect of the LFIB correction may be due to the MDT employed and to the different sources of the atmospheric forcing - the observational correction is derived from an ECMWF analysis, while the atmospheric forcing in FOAM is derived from Met Office NWP analyses. To make the most consistent use of these observations, atmospheric corrections should be derived from Met Office atmospheric models and the mean profile used to derive the SLA from the SSH should be produced from SSH measurements with only the desired corrections applied.

4.3 Filtering tides from model output

The tidal signature is the largest amplitude signal in the observations and models. Although in principle we could make use of this information, in practice we do not have an extensive enough time-series to improve the tidal prediction and the true tidal frequencies could be aliased by the observational sampling rate. We therefore wish to remove the tidal signature from both the observations and models. For the TAPAS observations, a tidal correction is supplied which is produced from a global ocean tide model. However, for the FOAM system we do not have separate output fields detailing the tidal elevation.

In the previous section we removed tides from the model using a Doodson filter. Since the actual tides are a complex superposition of many periods, any filtering algorithm is necessarily approximate. Although more complex tidal filters are available, Pugh [1996] found little difference in the calculation of monthly mean sea level with a Doodson filter or averages over longer time-spans.

Here we have repeated the comparisons for the case of no tidal suppression, and where we have removed the model tidal signature with a simple 25 hour running average instead of a Doodson filter. Figure 8 demonstrates that the tidal filtering affects only the lowest frequency variations along each example track, but the results in Fig. 9 show that the differences are minor. Overall the lowest standard deviation is achieved when the tidal signal in the model field is suppressed using the Doodson filter. However, since the residuals in the observation-model comparisons are not significantly affected by the filtering of tides, or indeed the choice of tidal filtering algorithm, we recommend that a simple 25 hour average is the best method to use operationally.

4.4 Spatial power spectra

Before removal of the large-amplitude tidal signature, the observed and modelled SLA are highly correlated. After filtering of the tidal signature, the correlations decrease suggesting that the next largest amplitude signal observed is not as well-matched by the model. From Fig. 2, it is clear that the highest frequency along-track variations seen in the observations are not present in the model, but that there are some features of the observations that are well-represented (e.g., the $\sim$200 km wavelength features seen in the bottom-left panel of Fig. 2), even if overall the correlation is low.

Since our hope is that there is some information in the observations that could be usefully assimilated, we have compared their power spectra with those of the models. Figure 10 shows the

© Crown Copyright 2013

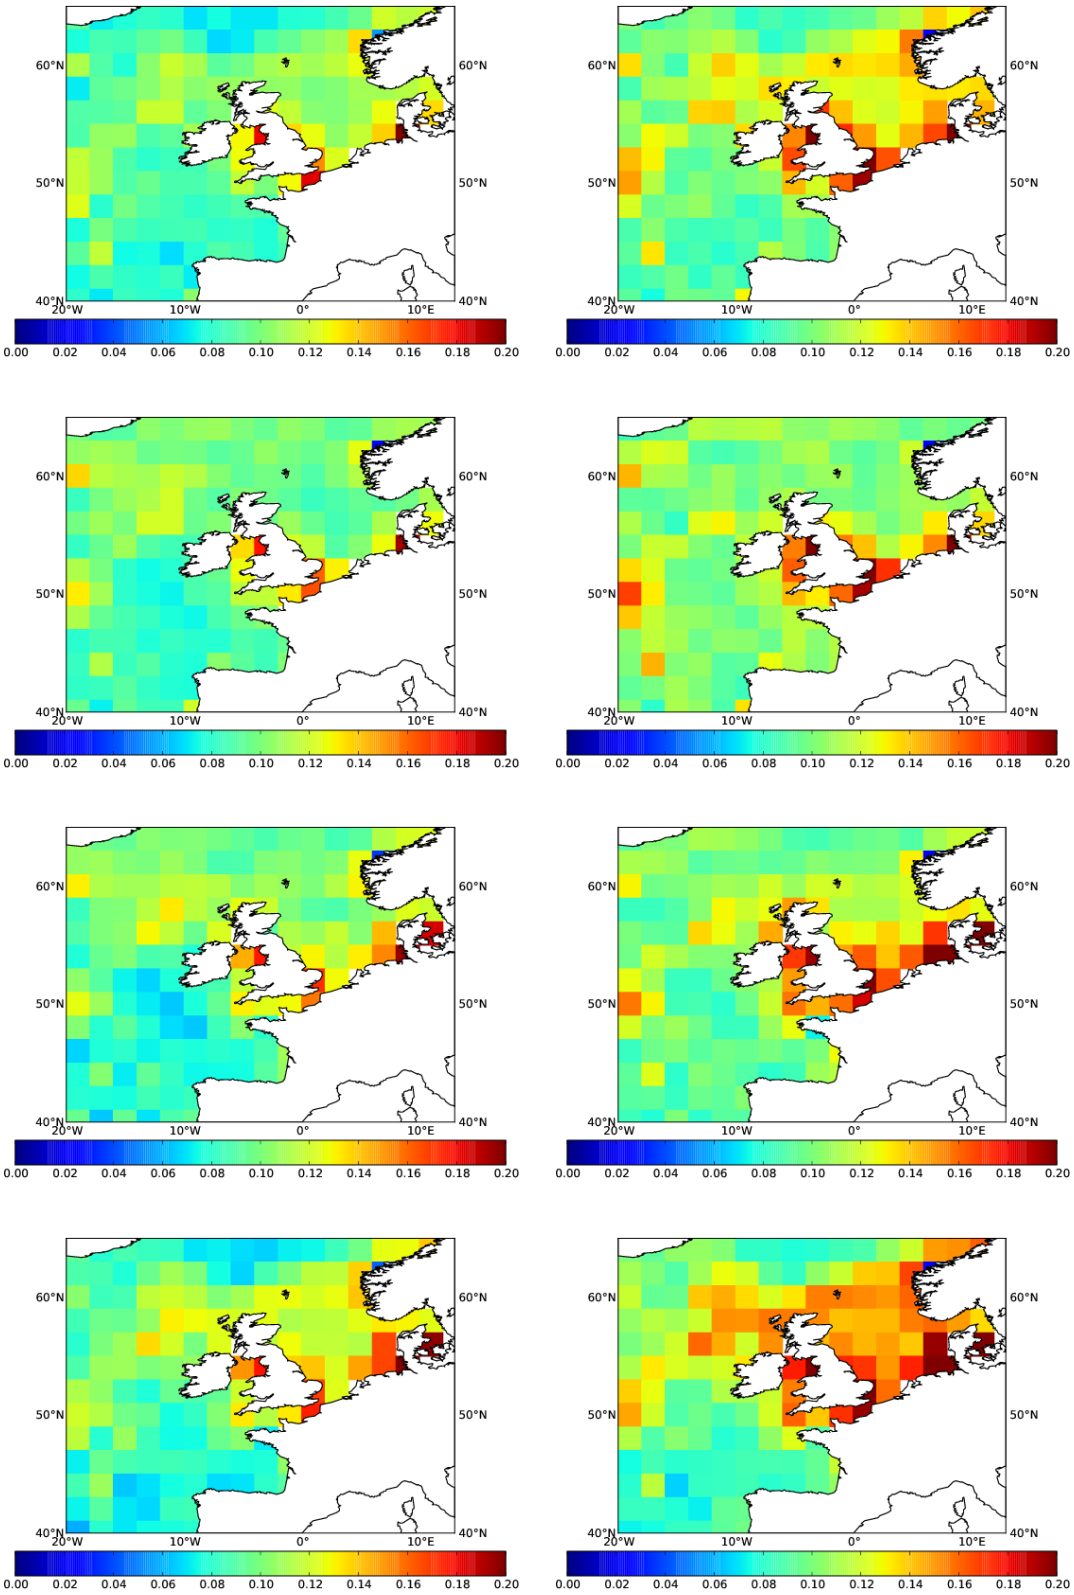

Figure 3: The absolute mean difference (left column) and standard deviation of the difference (right column) between the model and observed SLA from Jason-1 (for 01 Jan. to 19 Oct. 2008) with the following corrections applied (top to bottom): all except the MOG2D correction, all except the MOG2D and LFIB corrections, all except the LFIB correction, and all corrections applied. Here a Doodson filter has been applied to the model values to suppress the tidal constituents.

© Crown Copyright 2013

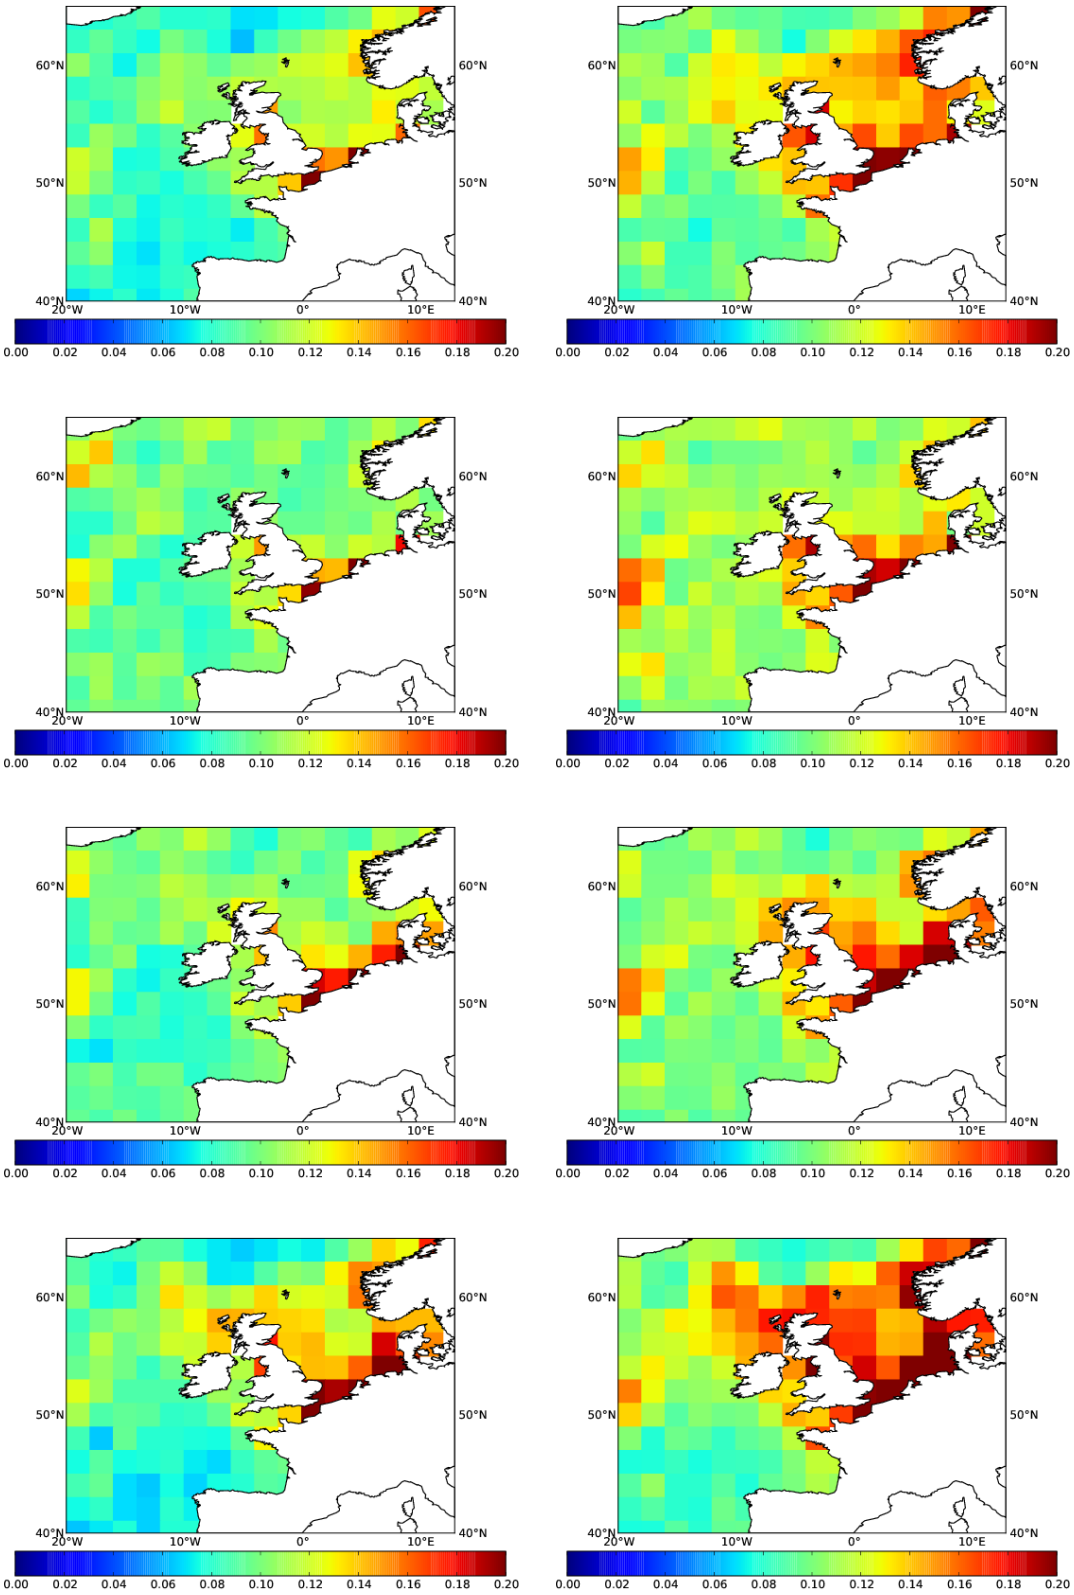

Figure 4: The absolute mean difference (left column) and standard deviation of the difference (right column) between the model and observed SLA from ENVISAT (for all of 2008) with the following corrections applied (top to bottom): all except the MOG2D correction, all except the MOG2D and LFIB corrections, all except the LFIB correction, and all corrections applied. Here a Doodson filter has been applied to the model values to suppress the tidal constituents.

© Crown Copyright 2013

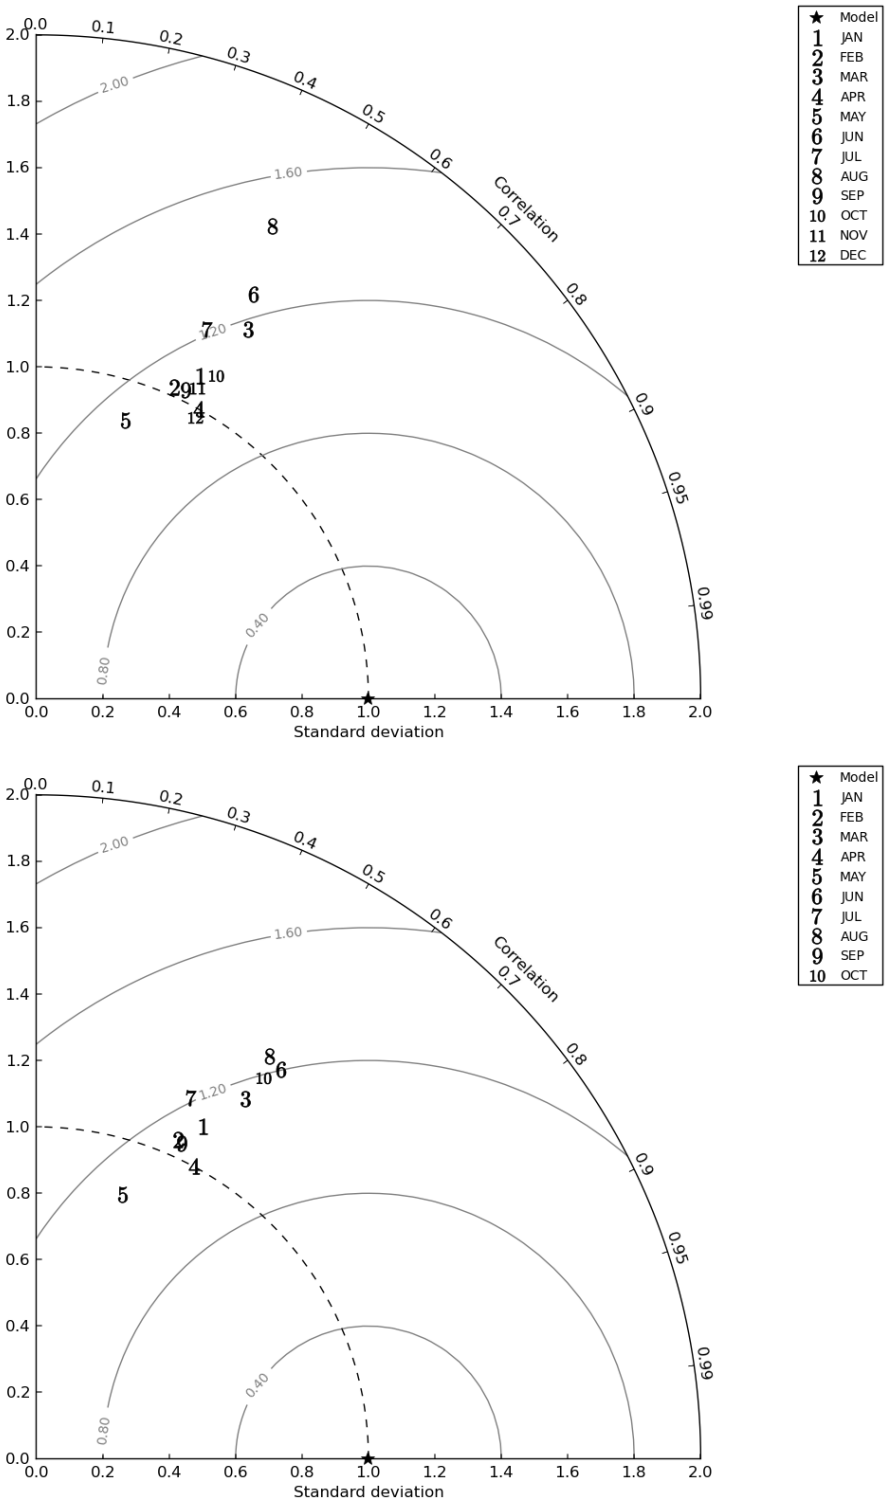

Figure 5: Taylor diagrams for ENVISAT (top) and Jason-1 (bottom) observations where no MOG2D correction has been applied compared to models where tides have been suppressed with a Doodson filter. Each panel shows the monthly observation-model comparisons. Since the model standard deviation varies from month to month, the radial axis shows the ratio of the standard deviations of the observations and model.

© Crown Copyright 2013

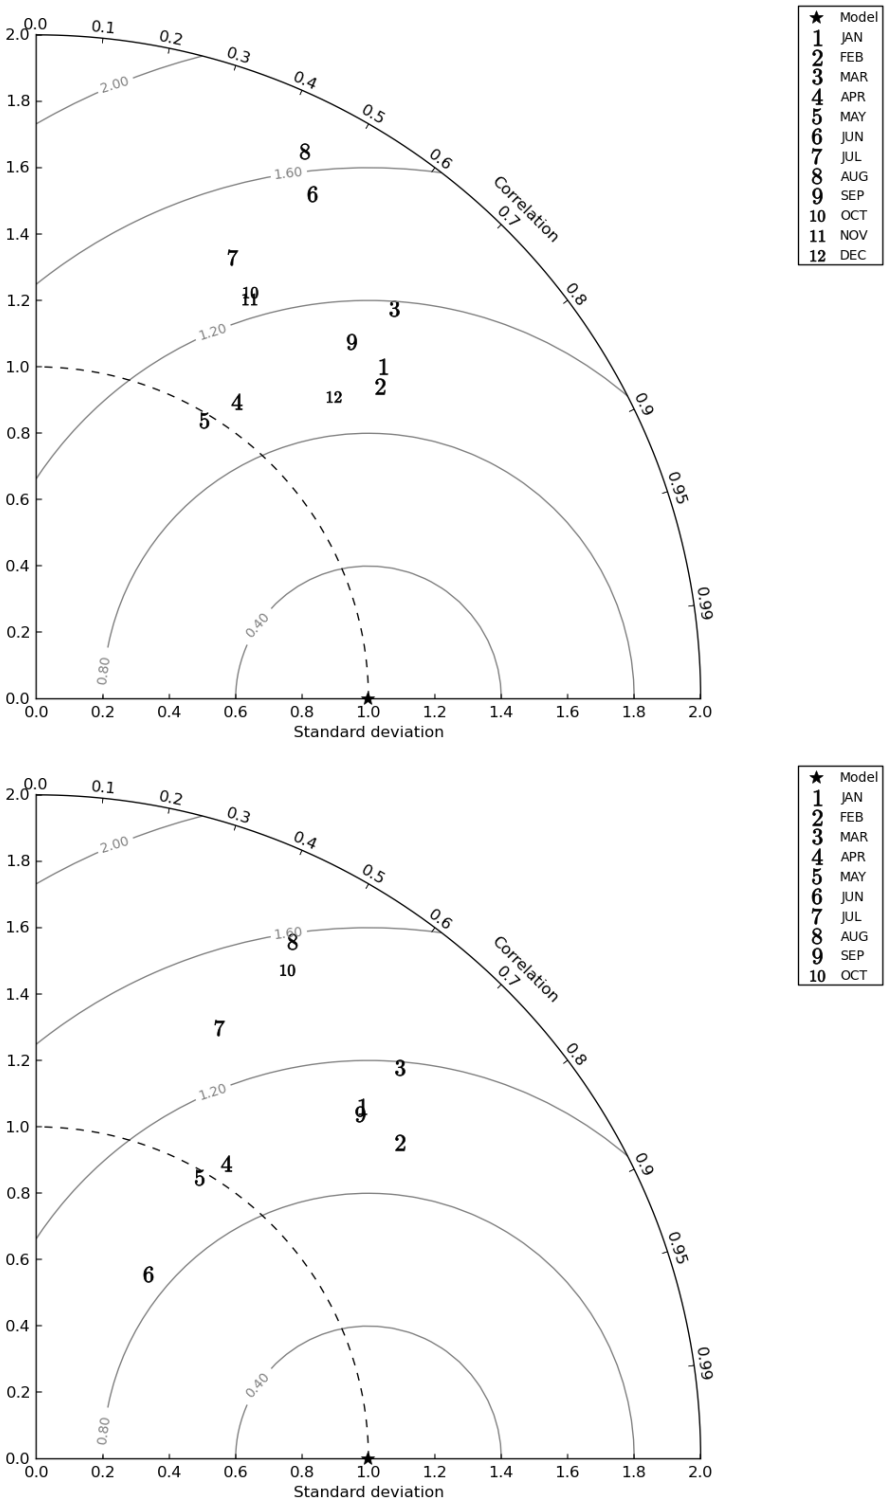

Figure 6: Taylor diagrams for ENVISAT (top) and Jason-1 (bottom) observations where no MOG2D or LFIB correction has been applied compared to models where tides have been suppressed with a Doodson filter. Each panel shows the monthly observation-model comparisons. Since the model standard deviation varies from month to month, the radial axis shows the ratio of the standard deviations of the observations and model.

© Crown Copyright 2013

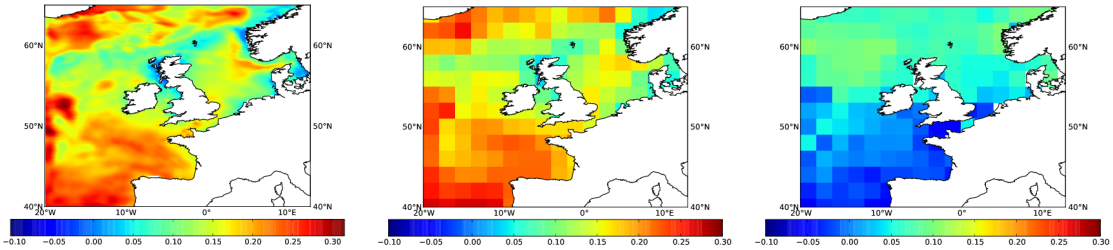

Figure 7: The difference between the Rio et al. [2011] mean dynamic topography (MDT) and the 1-year mean model field used to determine the model SLA (left), along with a version of the same gridded onto the same resolution as the observation-model statistics (middle). The right-hand panel shows the mean observation-model differences (for ENVISAT) where neither of the atmospheric corrections have been applied. The different MDTs used may be the cause of the structure seen in the right-hand plot.

power spectra of the observations and the associated model values for the individual example tracks discussed previously. These show that there is significantly less power at high frequencies in the model than seen in the observations. This is likely to be due to a combination of actual physical processes affecting the ocean state which are not included in the model, and due to observational noise.

Since the power spectrum of an individual track is relatively noisy and only gives a view of the observational-model differences over a thin track of ocean at one time, in Fig. 11 we have aggregated the information on the individual peaks in each power spectrum for every track over the year of observations. At each wavelength sampled by the observations, this shows the number of peaks with significant power at that wavelength and the mean power of those peaks. The minimum power limit of 10^{-2} m^2/km was chosen to exclude the many low signal-to-noise peaks present. Only Jason-1 data are used in the comparison in Fig. 11, but the differences relative to a comparison with ENVISAT observations are negligible. Since we expect the shelf to have more complex dynamics than the deeper ocean, we have carried out this analysis over the full region, and also with on- and off-shelf regions.

From Fig. 11, we can see that in all regions there is a large discrepancy between the number of significant-power peaks seen in the observational power spectra and in the model power spectra. This is consistent with the lack of high frequency model variations apparent in Figs. 2 and 10. In the off-shelf region, the removal of the tidal signature has little effect on either the number or the mean power of peaks in the observational power spectra. There are also very few peaks at high frequency in the model power spectra. On the other hand, in the on-shelf region, the tidal filtering removes some of the power from the observational power spectra at wavelengths below $\sim$100 km and almost all power below $\sim$100 km from the model.

The agreement between observations and model at wavelengths above $\sim$100 km in both on- and off-shelf regions, coincides with the expected spatial frequency below which observational noise dominates the signal [Fu et al., 2010] and there is a better correlation between the observations

© Crown Copyright 2013

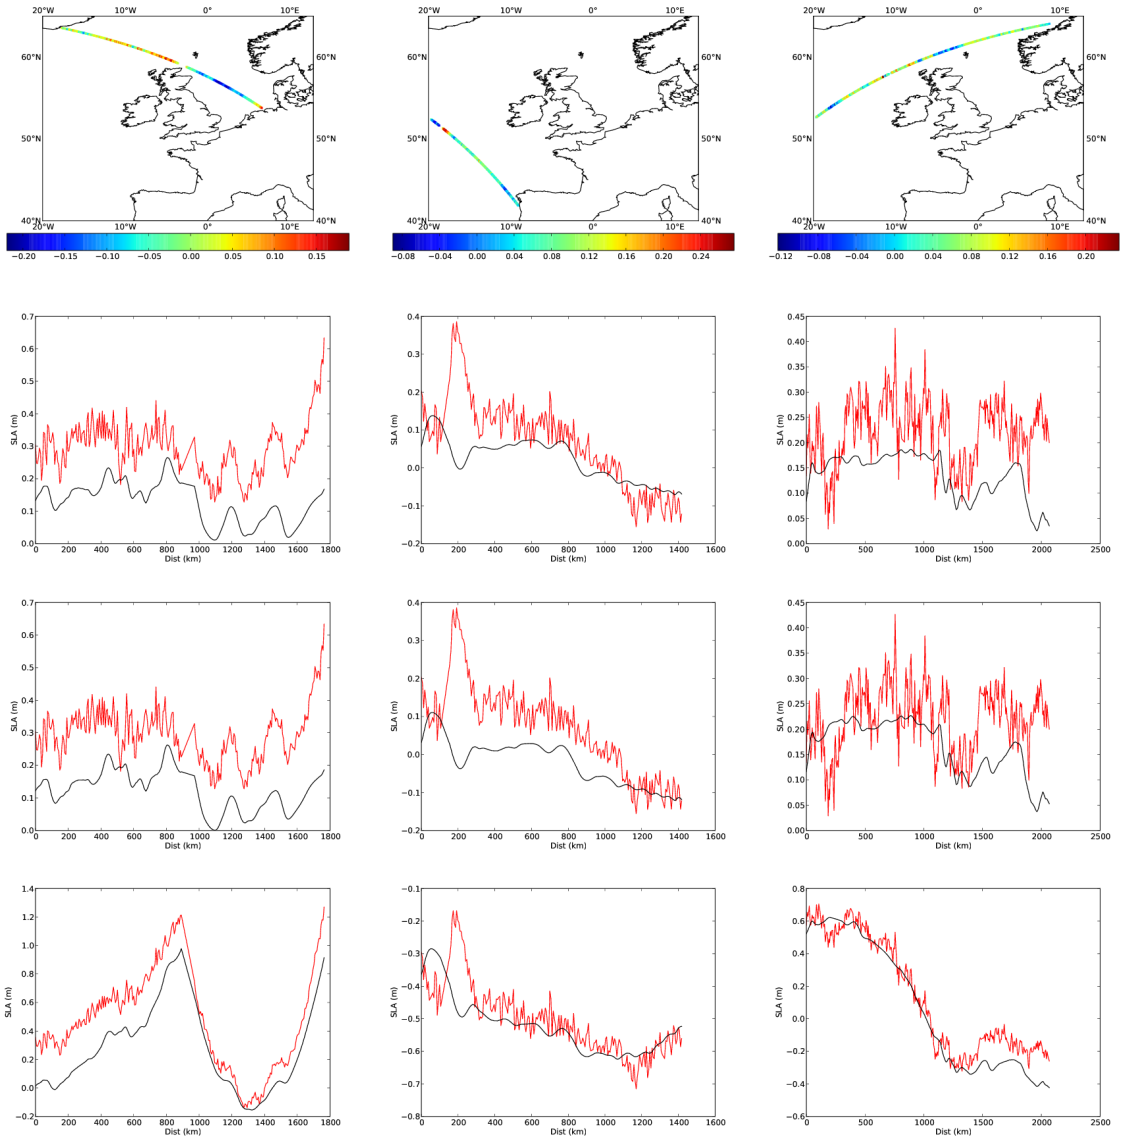

Figure 8: The effect of different tidal filtering on three individual tracks (left-to-right). The path of each track is shown in the maps on the first row and the subsequent rows show the model prediction (black line) compared to the observed values where neither of the atmospheric corrections (MOG2D and LFIB) have been applied (red line). In the second row plots, tides have been removed from the model using a Doodson filter, on the third row tides have been removed using a 25-hour running average, and on the bottom row tides are present in both the observations and model. Note the larger vertical scale where tides are present.

© Crown Copyright 2013

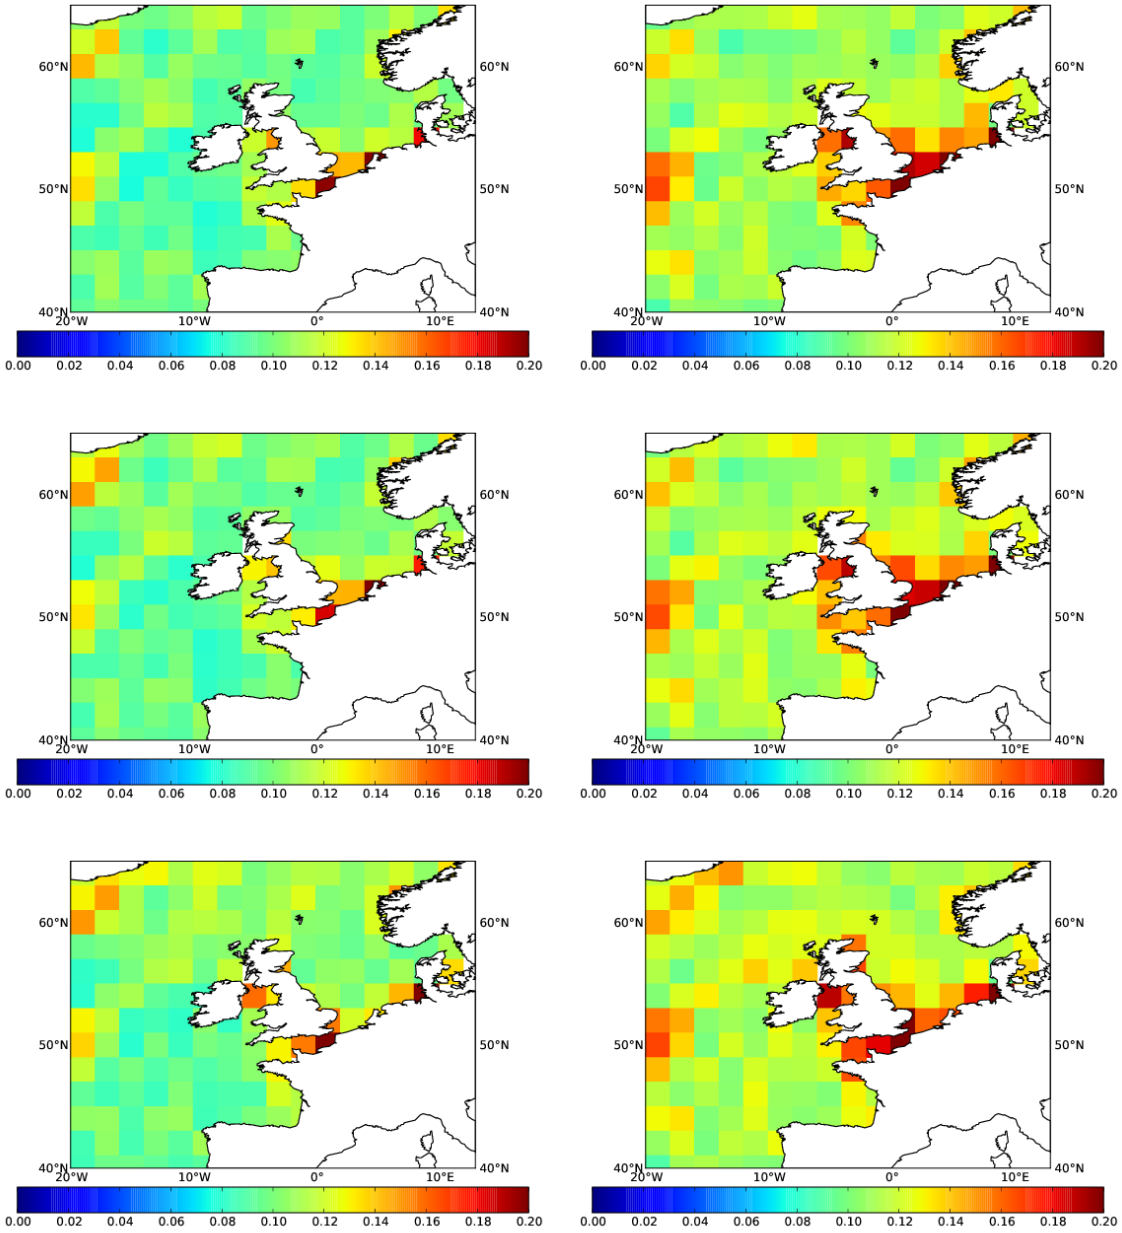

Figure 9: The effect of different tidal filtering on the FOAM-TAPAS comparison. The absolute mean difference (left column) and standard deviation of the difference (right column) between the model and observed SLA from ENVISAT (for all of 2008) with neither the MOG2D nor LFIB correction applied. The top row has had a Doodson filter applied to the model values to suppress the tidal constituents, the middle row has had tides suppressed by taking a 25-hour running mean, and the bottom row has had no tidal filtering from the model or observations.

© Crown Copyright 2013

and model at wavelengths greater than ∼100 km. Since the model does not contain the high frequency variations seen in the observations (whether real or not), these could be filtered from the observations before assimilation of the lower frequency information.

5 CS3X Surge Model Comparison

The FOAM system uses the Nucleus for European Modelling of the Ocean (NEMO) dynamical core which allows a single ocean model to be employed on both global and regional scales, and from short-term forecasting to climate prediction. However, NEMO was not developed as a shelf-sea model and so lacked some shelf-specific features [see O'Dea et al., 2012]. On the other hand, the CS3X surge model is a mature, well-tested and verified model tailored to the North-West Shelf and is the operational ocean model for surge forecasts over the NWS at the Met Office. FOAM has been running operationally for several years and so it is reasonable now to compare these new observational data against both models to judge their relative merits.

The CS3X surge model covers the NWS with a horizontal resolution of ∼12 km. It is forced by the sea-level pressure and 10 m wind fields and includes tides. The model is run in a fully-forced and a tide-only mode to allow the tidal and storm surge components to be separated. The surge model is run every 6 hours to produce a 48 hour forecast with model output (including surge and tidal elevations) every 15 minutes. To allow a comparison with the TAPAS data and with the FOAM model, we have combined the initial 6 hours of each forecast to create an unbroken, though necessarily discontinuous, forecast covering all of 2008.

Figures 12 and 13 show the absolute mean difference and the standard deviation of the difference between the observations (without either of the atmospheric corrections applied) and the surge elevation from the CS3X surge model (for both ENVISAT and Jason-1) along with FOAM comparison. Figures 14 and 15 show a similar comparison, but for the surge and tide elevation from the CS3X surge model compared to observations where the tidal correction has not been applied. These figures demonstrate that the CS3X surge model is more consistent with the satellite altimetry observations. While both FOAM and the surge model have higher residuals in specific areas such as the English channel and along the coast from France to Denmark, in general the CS3X surge model residuals are approximately 70% of those for FOAM. This is the case whether or not tides are present in the model and observations. However, there is a small increase (1–2 cm) in both statistics when tides are present in the observations and model. This may be due to a combination of the limited accuracy of the tidal models and the low cadence of the observations relative to the variations in the tidal elevation. Since the CS3X surge model fields are available at 15 minute intervals, compared to the hourly FOAM fields, a comparison was also made using only the hourly instantaneous CS3X surge elevation. This found only negligible differences due to the temporal sampling of the model fields.

Although there is a marked difference in the residuals between the two models, Fig. 16 shows

© Crown Copyright 2013

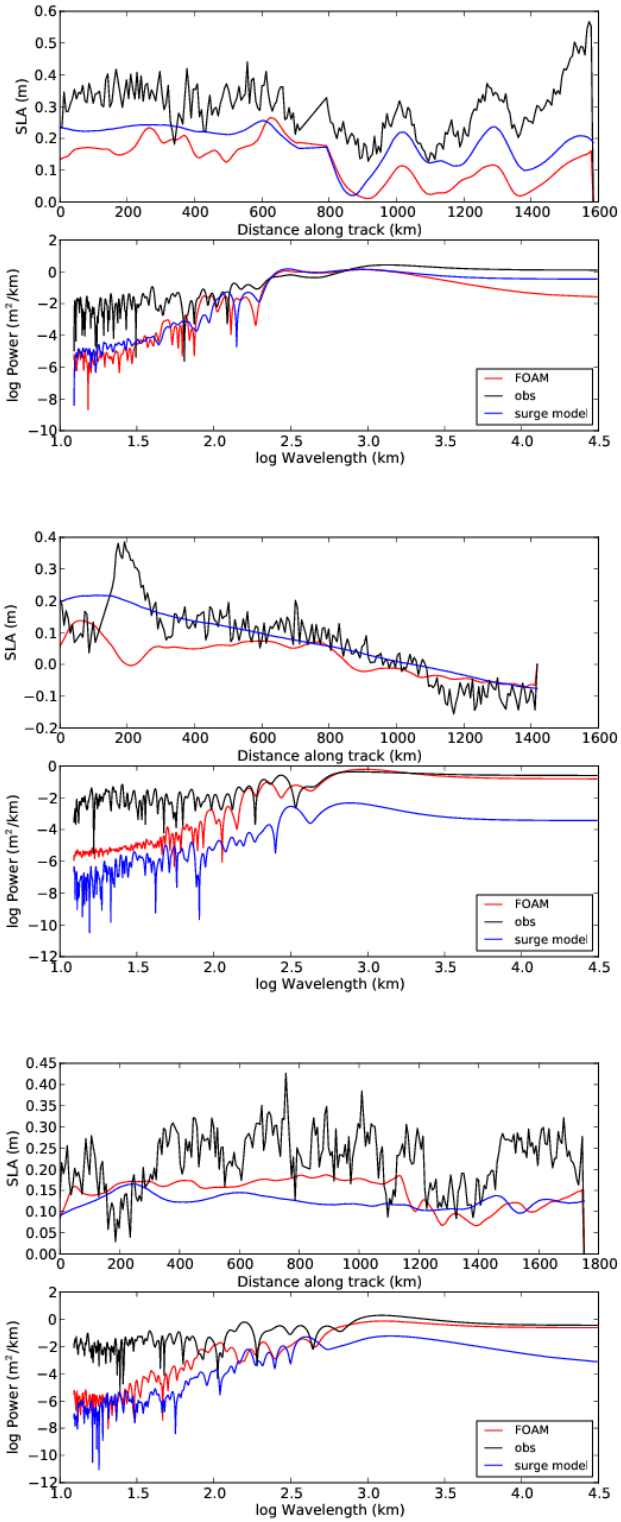

Figure 10: Single-track observation-model comparisons and power spectral density plots for each of the three example Jason-1 tracks (black lines) with the associated FOAM (red) and CS3X surge model comparisons (blue). The observations have not had the MOG2D or LFIB corrections applied and the tidal signatures have been removed from both observations and models.

© Crown Copyright 2013

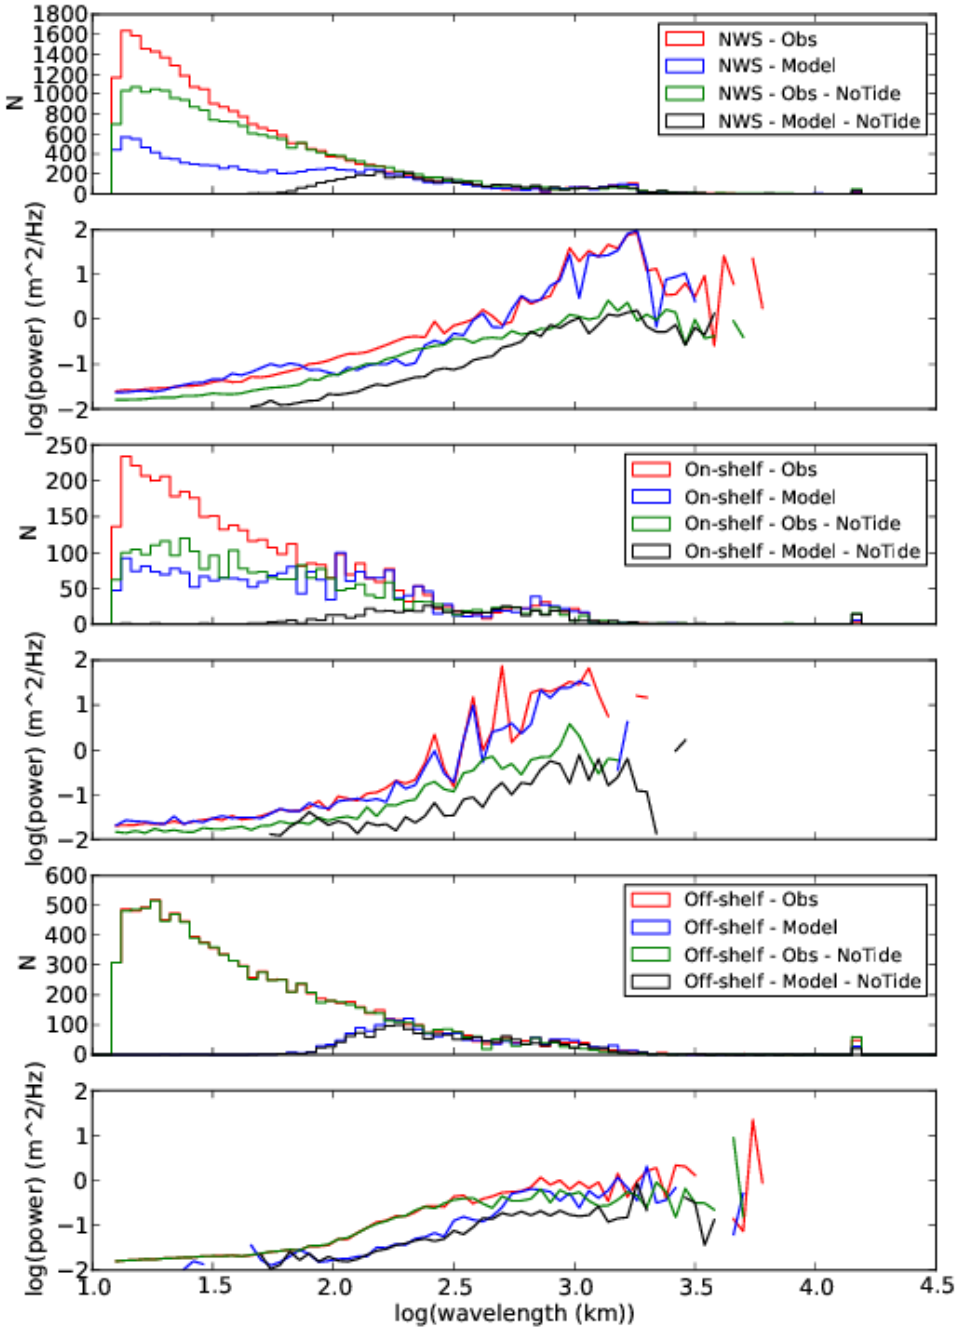

Figure 11: Histograms showing the wavelengths of all peaks in the power spectra of the Jason-1 observations and FOAM model (top panel) along with the mean power in each bin as a function of wavelength (second top panel). The three pairs of panels show the full NWS region, an on-shelf ($-5 < \lambda < 10$, $50 < \phi < 60$) and an off-shelf region ($-20 < \lambda < -10$, $40 < \phi < 65$). Each panel includes model and observation information, with and without tides. Cuts have been applied to the power spectra to ignore any with fewer than 100 observations or a mean observation spacing >10 km and to exclude peaks which have power $< 0.01\, m^2/Hz$.

© Crown Copyright 2013

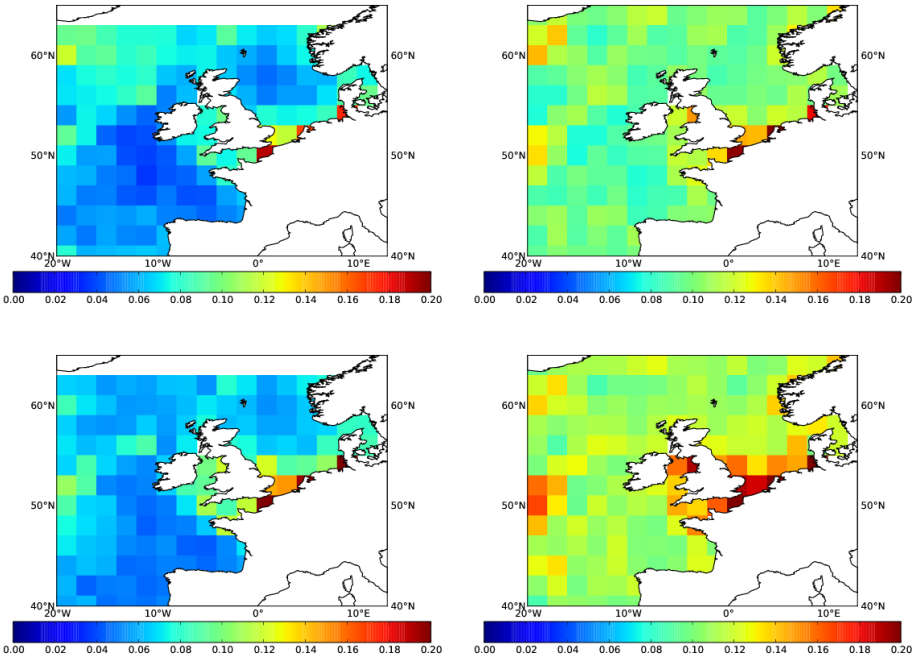

Figure 12: Absolute mean difference (top row) and standard deviation of the difference (bottom row) between the ENVISAT observations (with all except the MOG2D and LFIB corrections applied) and the CS3X surge level model (left) and the FOAM model with tides filtered using a Doodson filter (right).

that the power spectra are very similar and the removal of tides results in a suppression of wavelengths below ∼100 km as in the FOAM comparison. The comparison of the power spectra for three example tracks in Fig. 10 shows that the CS3X surge model is somewhat smoother than FOAM which may indicate that it contains less power at higher frequencies.

6 Summary and Discussion

In Sec. 4.1, we discussed how the tidal correction supplied with the TAPAS observations could be used to suppress the largest signal in the SLA observations. Since global ocean tide models assimilate many years of altimeter observations to predict the tidal elevation, and the low cadence of the observations may alias the tidal frequencies, it seems sensible to apply the tidal correction to exclude this signal from any assimilation of the SLA observations. The long wavelength correction supplied to minimise the effect of imprecise orbital information is small, but of presumed benefit.

However, the comparison of differences between Jason-1 and ENVISAT observations and the model predictions from FOAM have demonstrated that the atmospheric corrections (MOG2D and LFIB) should not be applied when comparing to shelf seas models. Since FOAM includes forcing due to atmospheric pressure and wind speed, it is sensible not to apply these atmospheric corrections to the observations. If either of the atmospheric corrections were used, it would be more

© Crown Copyright 2013

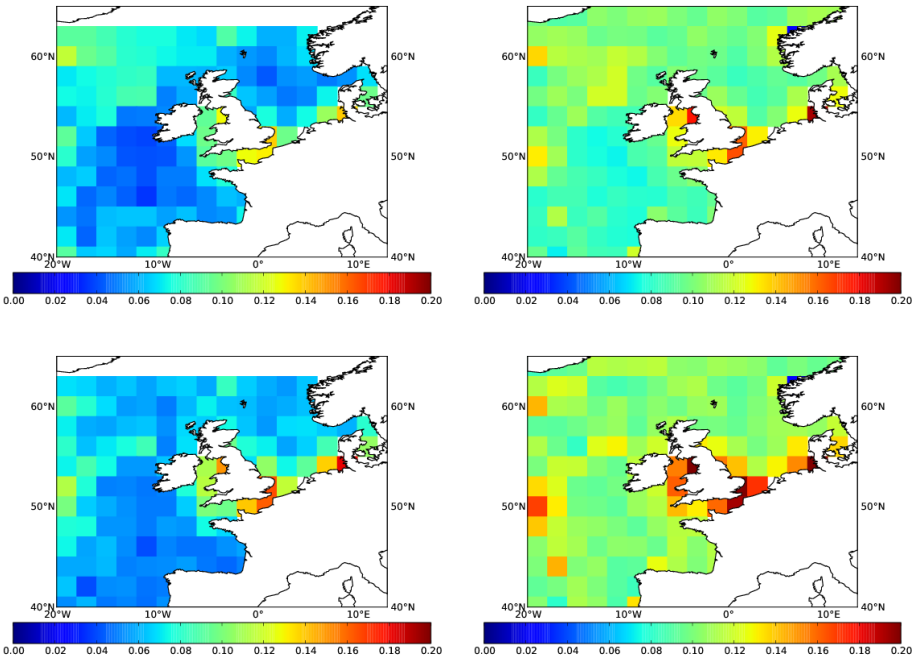

Figure 13: Absolute mean difference (top row) and standard deviation of the difference (bottom row) between the Jason-1 observations (with all except the MOG2D and LFIB corrections applied) and the CS3X surge level model (left) and the FOAM model with tides filtered using a Doodson filter (right).

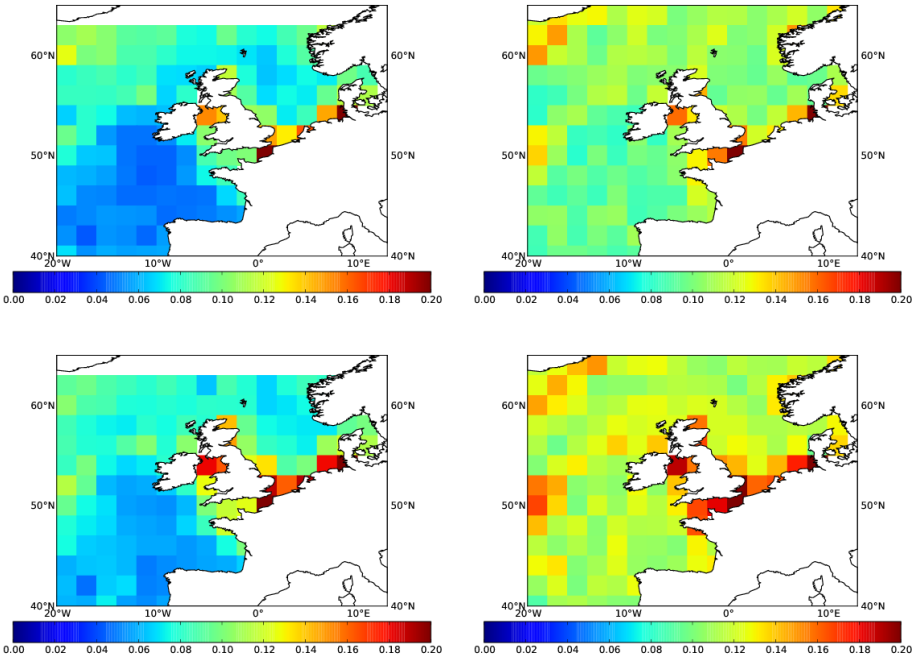

Figure 14: Absolute mean difference (top row) and standard deviation of the difference (bottom row) between the ENVISAT observations (with no tidal, MOG2D or LFIB corrections applied) and the CS3X surge and tide level model (left) and the FOAM model (right), i.e the same comparison as Fig. 12, but with tides included.

© Crown Copyright 2013

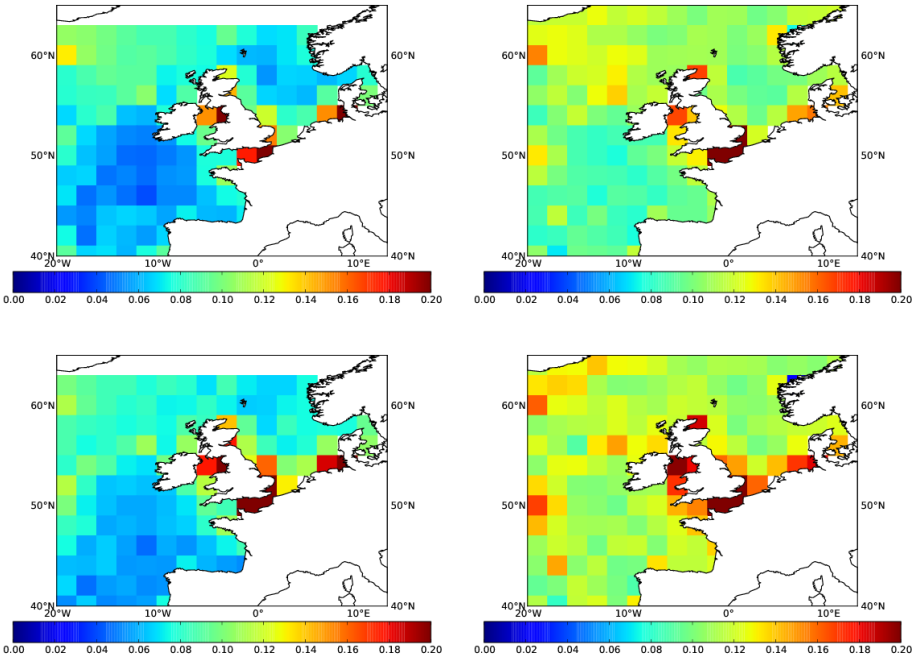

Figure 15: Absolute mean difference (top row) and standard deviation of the difference (bottom row) between the Jason-1 observations (with no tidal, MOG2D or LFIB corrections applied) and the CS3X surge and tide level model (left) and the FOAM model (right), i.e the same comparison as Fig. 13, but with tides included.

consistent if they were derived from Met Office atmospheric analyses rather than from an ECMWF analysis. In Sec. 4.3, we have shown the effects of the TAPAS-FOAM comparison when tides are not filtered from either model or data and when tides are filtered by means of a simple 25-hour running average, or using a Doodson filter. While the Doodson filter is superior to a running average, the benefit is small.

The maps of the absolute mean difference and the standard deviation of the difference between the FOAM predictions and altimeter observations (with only the LWE and tidal corrections applied) are shown in Figs. 3 and 4. From these we can see that over the full NWS region there is an absolute mean difference over the year of $\sim$10 cm, with up to 20 cm in the English channel and $\sim$8 cm off the South-West coast of Ireland. The values are marginally higher for ENVISAT than for Jason-1. The standard deviation of the differences over the year is also between $\sim$12 cm (Jason-1) and $\sim$13 cm (ENVISAT) increasing to $\sim$20 cm in the English channel for Jason-1 and over a wider area of the southern North Sea for ENVISAT. Instrumental uncertainties for the two satellites [$\sim$2 cm for Jason-1 and $<$4.5 cm for ENVISAT, Peng and Wu, 2009, Sunda, 2009] may account for the slightly lower differences in the Jason-1 comparison. Figure 7 also shows the structure present in the observation-model mean difference when neither of the atmospheric corrections are applied and contrasts this with the difference between the MDT used to construct the model SLA and the Rio et al. [2011] MDT. This suggests that the large-scale structure in the residuals may be due to

© Crown Copyright 2013

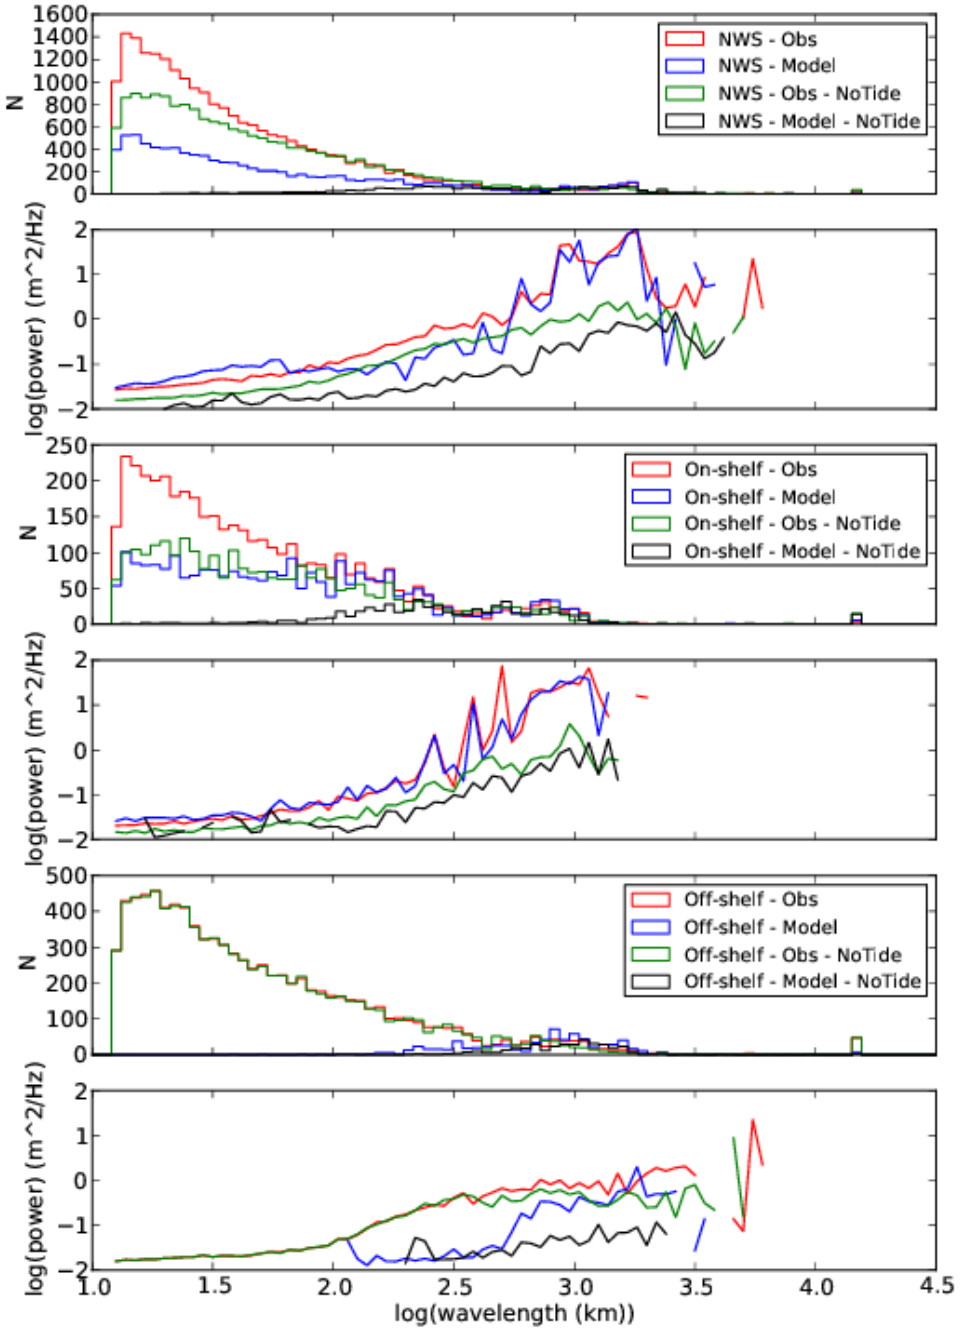

Figure 16: Histograms showing the wavelengths of all peaks in the power spectra of the Jason-1 observations and CS3X surge and tide model (top panel) along with the mean power in each bin as a function of wavelength (second top panel). The three pairs of panels show the full NWS region, an on-shelf ($-5 < \lambda < 10$, $50 < \phi < 60$) and an off-shelf region ($-20 < \lambda < -10$, $40 < \phi < 65$). Each panel includes model and observation information, with and without tides. Cuts have been applied to the power spectra to ignore any with fewer than 100 observations or a mean observation spacing >10 km and to exclude peaks which have power $< 0.01\, m^2/Hz$.

© Crown Copyright 2013

the specific MDT used.

Figures 5 and 6 demonstrate how the standard deviations, and correlation between the observations and model, change from month to month. While it is not completely clear that one set of corrections (excluding the MOG2D correction, or excluding both MOG2D and LFIB) is superior, the maps show an improvement over the North Sea when neither of the atmospheric corrections are used.

The power spectra information shown in Figures 10 and 11 also demonstrates the large difference between the observations and model, especially at high frequencies. The tidal correction has little effect on either the model or observational power spectra in the off-shelf regions, but in the on-shelf regions the tidal correction suppresses the high-frequency power in the observational power spectra and almost removes it in the model power spectra.

In both on- and off-shelf regions, there is much better agreement between the observations and models at wavelengths above $\sim$100 km. We know that below this wavelength observational uncertainties can dominate the signal, and it also seems that the current model has little mesoscale variability (i.e., $\lambda \lesssim 100\,km$), but the surge component of the variability could be usefully assimilated. This is probably due to the relatively low resolution of the model compared to the baroclinic Rossby radius on the shelf. This 100 km boundary, above which physical variations included in the ocean model can be seen in the observations, provides a natural cut-off for frequencies which can be usefully assimilated. Since the model does not show the high frequency variations seen in the observations, by assimilating them we risk destabilising the model. On the other hand, if these variations are simply a result of the instrumental noise, then with a proper understanding of the errors we can assimilate all of the altimeter observations without pre-filtering of high frequencies.

The differences between the model-observation comparisons for FOAM and the CS3X surge model shown in Figs. 12 to 15 demonstrate that the predictions from FOAM are not yet as accurate as those from the more mature model tailored to the North-West Shelf. However, they have broadly similar power spectra and a detailed comparison of the two models will soon begin to better understand how FOAM can be improved.

Finally, aside from the differing physical and temporal scales of the dynamics of the ocean compared to what can be observed with current altimeters, an important difficulty with assimilating altimeter observations in the shelf seas relates to the seasonal stratification. In the deeper ocean, sea level anomaly observations can be assimilated by introducing incremental changes to the temperature and salinity profiles of the water column beneath the observation locations. Effectively, the depth of the thermocline is altered. However, the NWS is not stratified year-round and the sea level variability on the shelf is also heavily influenced by storms. A more physically realistic route to assimilate the observations may involve incrementing the surface elevation directly which would not depend on knowing about the stratification. Both of these methods are already possible in FOAM/NEMOVAR and so tests of altimeter assimilation in the NWS could attempt to use both methods to assimilate the information in the TAPAS data.

© Crown Copyright 2013

7 Conclusions and Recommendations

We have compared an experimental SLA product (TAPAS) with the predicted SLA from FOAM and the CS3X surge model. By comparing the model fields and observations, with various corrections applied or excluded, we have concluded that the tidal and long wavelength error corrections should be applied, but neither of the atmospheric corrections supplied with the TAPAS observations should be used when comparing to FOAM.

FOAM includes a tidal component which is not known independently, so the model fields must be filtered before any comparison can be made. Also, since there is only a marginal improvement between using a 25-hour running average and a Doodson filter to remove the tidal signature, we recommend using a 25-hour average as this is significantly easier to implement operationally.

The TAPAS-FOAM comparison has shown that, averaged over a year, there is a $\sim$12 cm standard deviation and $\sim$10 cm absolute difference between the SLA observations and FOAM predictions. A similar comparison between the satellite altimetry and the CS3X surge model demonstrated that the FOAM predictions for the SLA in the NWS are not yet as accurate as the surge model, with residuals approximately 70% of those seen with FOAM.

From the spectral analysis we have shown that the mesoscale variability is not well-represented at the current resolution of FOAM and is dominated by instrumental noise in the observations. This suggests that, before assimilation, the observations may have to be low-pass filtered to remove variability that might introduce increments which would destabilise the model. The tidal component is also not well temporally sampled by the observations, but the surge component of the SLA variability is present in both observations and models. Since this component is barotropic, the SLA can be directly assimilated rather than relying on temperature and salinity increments. Future increases in the model resolution and spatial and temporal sampling of the observations may also allow the mesoscale component of the SLA variability to be assimilated into shelf seas models.

Even though there is a large difference between the spatial and temporal scales of the dynamics in the shelf seas and that observable by existing altimeters, these experimental altimeter products present an opportunity to utilise a substantial number of observations in a relatively data-starved, yet hugely important region. The level of improvement possible with the assimilation of these observations will become clear during assimilation tests, but any improvements in the European North-West Shelf will translate to other shelf seas models, and the comparison made here demonstrates how the model fields and observations must be filtered and corrected before the observations are assimilated.

© Crown Copyright 2013

References

James A Cummings, Laurent Bertino, Pierre Brasscur, I Fukumori, M Kamachi, ML Martin, P Oke, CE Testut, Jacques Verron, and Anthony Weaver. Ocean data assimilation systems for godae. *Oceanography*, 22:3, 2009.

P De Mey, P Craig, F Davidson, CA Edwards, Y Ishikawa, JC Kindle, R Proctor, KR Thompson, and J Zhu. The godae coastal and shelf seas working group (csswg) community. 2009. applications in coastal modeling and forecasting. *Oceanography*, 22:198–205, 2009.

S Dobricic, C Dufau, P Oddo, N Pinardi, I Pujol, and M-H Rio. Assimilation of sla along track observations in the mediterranean with an oceanographic model forced by atmospheric pressure. *Ocean Science*, 9:1577–1598, 2012.

GD Egbert and RD Ray. Significant dissipation of tidal energy in the deep ocean inferred from satellite altimeter data. *Nature*, 405:775–778, 2000.

Jonathan Flowerdew, Kevin Horsburgh, Chris Wilson, and Ken Mylne. Development and evaluation of an ensemble forecasting system for coastal storm surges. *Quarterly Journal of the Royal Meteorological Society*, 136:1444–1456, 2010.

L Fu, Douglas Alsdorf, Erensto Rodriguez, Rosemary Morrow, Nelly Mognard, Juliette Lambin, Parag Vaze, and Thierry Lafon. The surface water and ocean topography (swot) mission. *Proceedings of OceanObs*, 9, 2010.

Gurvan Madec. Nemo ocean engine. *Note du Pole de modelisation, IPSL*, pages 1288–1619, 1999.

MJ Martin, A Hines, and MJ Bell. Data assimilation in the foam operational short-range ocean forecasting system: a description of the scheme and its impact. *Quarterly Journal of the Royal Meteorological Society*, 133:981–995, 2007.

EJ O'Dea, AK Arnold, KP Edwards, R Furner, P Hyder, MJ Martin, JR Siddorn, D Storkey, J While, JT Holt, et al. An operational ocean forecast system incorporating nemo and sst data assimilation for the tidally driven european north-west shelf. *Journal of Operational Oceanography*, 5:3–17, 2012.

DongJu Peng and Bin Wu. Precise orbit determination for jason-1 satellite using on-board gps data with cm-level accuracy. *Chinese Science Bulletin*, 54:196–202, 2009.

N Pinardi and G Coppini. Operational oceanography in the mediterranean sea: the second stage of development. *Ocean Science*, 6:263–267, 2010.

David T Pugh. *Tides, surges and mean sea-level*. John Wiley & Sons Ltd., 1996.

Richard D Ray. A global ocean tide model from topex/poseidon altimetry: Got99. 2. *NASA Technical Memo*, 209478, 1999.

© Crown Copyright 2013

MH Rio, S Guinehut, and G Larnicol. New cnes-cls09 global mean dynamic topography computed from the combination of grace data, altimetry, and in situ measurements. *Journal of Geophysical Research*, 116:07018, 2011.

D Storkey, EW Blockley, R Furner, C Guiavarc'h, D Lea, MJ Martin, RM Barciela, A Hines, P Hyder, and JR Siddorn. Forecasting the ocean state using nemo: The new foam system. *Journal of Operational Oceanography*, 3:3–15, 2010.

M Sunda. Envisat-1 products specifications. *ESA Technical report*, PO-RS-MDA-GS-2009, 2009.

The CNES-CLS09 MDT was produced by CLS Space Oceanography Division and distributed by Aviso, with support from CNES and the TAPAS altimeter products were produced and distributed by Aviso (http://www.aviso.oceanobs.com/), as part of the Ssalto ground processing segment.

© Crown Copyright 2013

Met Office
FitzRoy Road, Exeter
Devon, EX1 3PB
UK

Tel: 0870 900 0100
Fax: 0870 900 5050
enquiries@metoffice.gov.uk
www.metoffice.gov.uk

Produced with the Met Office LaTeX template v2.0.0